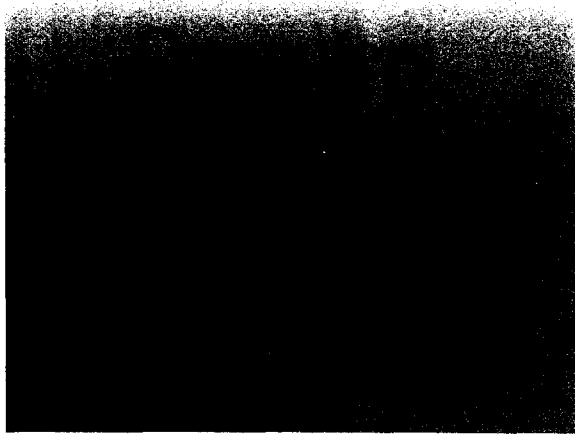


**Proceedings
of the Symposium
on Medical Informatics**

Sponsored
by the
Association of American Medical Colleges



In this decade, discussing the role of computers in medical education is like arguing for the role of books. Like books, the electronic *media* are just that — media for the expression of canonical knowledge in a form that is subject to organized criticism and review. In her masterful review of the cultural impact of printing, Elizabeth Eisenstein put less weight on the economy of multiple copies and their dissemination than on the faithful reproduction of an author's work enabling cumulative criticism and perfection of the text. At the time of Gutenberg's invention, the world already had millions of volumes of manuscript materials, and it took a century before the new products of the printing press overtook the accumulated manuscripts. The instant and critical change was that a particular work could be faithfully and accurately replicated, that it could be proofread, and that it could be subjected to criticism, and that the same text could be examined by a variety of readers. Authentic data and scientific information could now be recorded with incremental improvement, in place of a process which had produced new error and gloss and variation with each new copy.

Likewise, that critical and responsible examination of text by an author and by the scientific community is as important a function of electronic communication as is its rapidity and convenience of dissemination.

From my own particular experience, I am better qualified to talk about journals than about books, and about medical research than medical education. Dr. Myers did indicate that I have spent some decades in settings where medical students have been an important part of the environment. We ought to acknowledge the reality, nevertheless, that undergraduate medical education has taken second place in the actual preoccupation of most academic medical centers. Graduate medical education would be the more realistic focus of this discussion, as it is the senior educational responsibility at most of our institutions. And, Dr. Cooper has already pointed out the indispensable role of continuing medical education in a world where the medical knowledge base is changing so rapidly.

Computer literacy has the same role and ought to be taken for granted in much the same way as book literacy. In principle, students need not know

about electronic circuitry or software compilers any more than they need know about the mechanics of the printing press. A primary barrier to computer literacy, however, is fear about interfacing with a terminal, and this is perhaps more acute in the middle aged than in younger students. They must become sufficiently familiar with the terminal so that they are not overawed. The information that comes from the terminal is exactly as reliable as the information that was put into it, just as the information in the daily press or the printed volume deserves no more credit than is due its sources, with the exception that it is exposed to the critical judgment of the community.

One reason we can rely on some print media is that there is a better than evanescent record. Letters to the editor can provide feedback to a journal. Books may be issued in new and revised editions which should be more trustworthy. A vicious aspect of the broadcast media, radio and television, is the near impossibility for such feedback outside the courts. There is virtually no opportunity for cumulative correction and perfection in the small steps that we have in the other media.

Students should then be literate enough that they are not overawed, and so that they confidently develop an informed and critical attitude about the wisdom they receive at the terminal, just as they should for the information they receive in the lecture room or from the book. They need to understand the reliability of the systems, and above all, the extent to which these are ultimately totally based on human fallibility. The final argument for pedagogic investment in this area has less to do with computers *per se* than with the understanding of models of complex systems, their intricacy and principles of design, and how they should be critically reviewed. System models and software programs are ideal learning frameworks for sharpening intellectual skills and for providing experimentally testable perspectives that can be helpful in understanding other complex systems like human physiology.

The question of computer literacy is, however, pretty well moot now for most of our youngsters. With the universal penetration of personal computers into businesses, colleges, and high schools, we have a generation where the typical medical student is likely to be more literate in that domain than the typical full professor. Many youths now learn to spell through correcting their type-script in a text-processing machine, which will indicate where their words deviate from the dictionary provided. Hence, we may get a restandardization of language. For all those reasons I described, these new media provide the basis for a new literacy. But it still may be a long time before the senior physician gets over the social stigma of being seen operating a keyboard. Perhaps voice entry can cope with barked commands as input to computers, or pattern recognition systems will surmount the ultimate challenge of deciphering handwritten prescriptions.

Besides the role of computers as books, there is the very important function of computers as instruments. In this capacity they are more and more embedded into daily medical practice. We have CT scanners, NMR scanners, and similar devices. As unit prices plummet, we will see more office-based instru-

mentation proliferating in the diagnostic enterprise, assisting the physician. Again, the physician does not have to know every detail of the electronic circuitry which controls the x-ray machine, but it is better if he does know something about it. This general knowledge is necessary not only for a sense of the limitations of the technology, but also so the user can identify new opportunities for innovation.

A particular book called the "patient chart" is rapidly being supplanted by computer-based hospital information systems. The management information system is more typical within hospital practice, and is being used for such activities as the integration of laboratory data, drug monitoring, adverse effects reports, and monitoring for drug interaction besides accounting and billing. We are beginning to see that expert systems can oversee the mechanics of the practice of medicine and provide some buffering against unintended errors. A sense of how these things operate will be indispensable to their efficient use by physicians.

As government regulations grow to control the economic flows involved in medical practice, Professional Review Organizations will hardly be able to function effectively, nor can payment based on diagnostic related groups be implemented, without a substantial degree of automation. This also implies substantial supervision by others in the practice of medicine.

In turning to some more general features, the use of computer technology can be divided into the learning process (in which the user is explicitly a student), the practice process, and the research process, which provides the fundamental knowledge and expertise required for learning and practice. It is important that these not be too far apart from one another. If they are not integrated within the same medium of communication, or if research, learning, and practice cannot be expressed in a common language, or if they do not rely on a consistent data base, then there is chaos, futility and disaster.

Computers are important factors in the research process, and computers in communications networks will have a role in the generation of scientific consensus in the research process. I need not elaborate further on the use of computers in embedded instruments, in the analysis of data, or for the conduct of statistical studies. Many labs have their laboratory notebooks on-line with direct acquisition of data from experiments. This is a mechanical or engineering function of those systems. The new usage will be for the critical interaction of experts with one another: in an explicit community which may be working together on a particular problem or in an invisible college of individuals who may be exchanging information with one another. An example of the latter is the BIONET community sponsored by the National Institutes of Health, which allows individuals working in the field of molecular genetics to share a common, and now vast, database. There are 3 or 4 million characters of DNA nucleotide information assembled from around the world. A center in Philadelphia collects information for protein sequences. Los Alamos has the contract for DNA data. There is an international collaboration covering the European sector. BIONET itself is a shared resource, centered in a mainframe system

in Menlo Park, California. The system not only provides access to those databases on sequence information for investigators throughout the world, but provides the software tools for doing all of the combinatorics, the matching of sequences, looking for the evolutionary patterns, and handling vast amounts of data. But again, the most important function of the media providing support to that community is the inter-communication among the individuals. BIONET itself is too young to have demonstrated that yet, but it was certainly the most exciting aspect of its progenitor, the SUMEX-AIM computer system at Stanford which helped introduce artificial intelligence methodology to medical research.

The learning uses of computers are those which most nearly approximate books. The student must have access to their consensus. I use consensus not with the meaning of a well-defined authoritative statement of present knowledge, but as a distillation and authentic exposition of the controversy persisting on a subject. Here we have the role of the database retrieval systems for bibliographic inquiries, the extension of MEDLINE which enables the student to discover what has been written. Although at the moment there is a limited amount of full text material in such systems, one can generally get at the abstract or at least the title and key words as pointers to the library books. We are rapidly reaching the point where there will be two tiers of records in the scientific literature. There will be the paper publications which will serve a function of gate-keeping indicators of recognition and prestige. The instant communication, or preprint level of dissemination, is already being taken over in many fields through sharing on computer-based communications media. Students must also be a part of this process because they cannot afford to be given out-of-date information on specialized problems. Students have as much right to up-to-date versions of knowledge as others in the research community, and learning how to participate in that process is essential preparation for being at the cutting edge of their future careers. In addition, books have mechanical limitations. They are static, linear, and may leave much to the imagination, which can be alternately desirable and frustrating. It is very hard to get a good feel for the three-dimensional aspect of DNA by looking at projections in a book; even a stereoscope gives you only part of the story. If you really want to understand the significance of the bond angles that have to be adjusted, and the location of the strains in the structure, or the strengths in the hydrophobic interaction of bases, you need something closer to a four-dimensional model, which includes the time axis. Here nothing can surpass computer graphics. The users' ability to interact with those graphics is gradually being transferred from being primarily a research tool into something students can follow and acquire some understanding about.

Much the same thing could be said about many physiological systems. I don't think one can really get a very good picture of the hemodynamic system without having a model that can present the sensitivity of the pressures and the flows to changes in the input variables. A student cannot test the authenticity of an author's statements about those models without having them in some opera-

ble, testable form. The same thing can be said for the structure of the central nervous system. One can probably get a far more authentic view of the gross morphology of the human body with the aid of computer-based representation of those three-dimensional structures than in those rather poor simulations that we have in the cadaver, which is such a poor model for living structures.

The computer can also provide a base in what is called computer-aided instruction or self-interrogation. Students can look at their mastery of the field in ways that are far more cost effective than having a professional tutor in constant attendance. I think we can all agree that the practice of intellectual skills in handling data is far more important than the futile effort to acquire all the knowledge which might be relevant to the practice of medicine at any given stage. The knowledge a physician needs can change overnight, but the physician's skills will have to persist. Acquiring skills such as adapting new information into one's framework of knowledge, to correlate this with the current art, and to check it for internal consistency can be done most effectively with the student's interaction with the knowledge base at the terminal. It is possible to do this now; there will still always be a place for interaction with live, warm bodies at some part of that process to provide text, to explain where things went wrong, or where the model didn't perform to expectations.

It is economically infeasible to offer intensive one-to-one relationships of professor to student which are needed for the full development of such metal-logical skills. I believe we will find that we need to use medical students with one another as peers as an important part of this process to get the intellectual talent that is needed to sharpen the critical function in the depth and detail which is required in the teaching process. There is nothing like having intellects with different world experiences to provide critical interaction under the ultimate supervision of a tutor. For this to work, we will need new kinds of figurative books, such as the expert systems now being developed. An example is the CADUCEUS system that Jack Myers is working on. This is an immense undertaking, with decade-persons of work needed to extract human knowledge, and to compose the abstractions from that knowledge into a coherent and mutually consistent logical system. This is the real challenge in the relationship of experts to these systems. For the first time we have machines that demand consistency; we have a kind of self-checking on our own knowledge. I remember the rude shocks I had about even quite elementary organic chemistry when I was starting to put in the rule systems for DENDRAL. We simply do not have in our own heads the machinery for detailed proof of the consistency of our knowledge bases. Cognitive inconsistencies emerge when all of that knowledge is put into one place and put through mechanized systems for consistency checking.

The practice side would not be too different. As it becomes obvious that it is impossible to convey more than a fraction of contemporary knowledge, we must have more emphasis on continuing medical education as a part of the physician's career. Since physicians in practice will be far more diverse in their experience and new needs than are medical students, they will have even

greater reliance on self-regulated, computer-aided learning systems.

The use of expert systems in medical practice may be tempered by patients' images of the profession. Will they accept the expertise of the information scientist—the one who best knows how to access the world's data—and pay the same fees as they do for the omniscience offered at first hand by the physician? Will they take more or less assurance from their doctor himself seeking “a second opinion?” Will doctors be able to sustain their self-assurance in such an environment, or will this be resisted in a struggle to sustain their traditional image?

This decade has given us new firm bridging between fundamental science in molecular biology, cell biology, and physiology and medicine. This pace of change forces us to abandon the notion that the function of undergraduate medical education is to transmit the wisdom of a given year. Undergraduate medical education has to be thought of as the means to provide the base for a life-long, continuing learning process. It will be impossible to manage that without the use of computer-based knowledge systems. At each of these levels, access to the book-based libraries of the world is an essential feature, as is their continuing replacement with more mechanized systems. At the core of the consensus building, in which we rationalize and make our knowledge coherent, are expert systems. During that same decade computer science concentrated in building the framework for expert systems, the computer software that could understand and apply rules. We now need new meta-systems to govern the logical structures of the expert systems themselves. But the hard part is yet to come, and that is in the extraction of the human expertise to be put into the systems. This is an unbelievably costly effort. If you have not yourself been the subject of a process which has tried to extract the knowledge you have in a given subject area, you are in for an interesting experience. The relentlessness of these systems, their demand that you say what you mean and that you mean what you say, and the extent to which rules that you impart are tested against all of the other instructions put into the system is quite a chastening experience for anyone, even in the field of their central expertise.

Right now expert systems may not be totally satisfactory. But, if you look at it really critically, no book is totally satisfactory either, and if you were to rely slavishly on the rules given in any single textbook, you would get the same disasters as if you were to rely slavishly on any now functioning expert system for exactly the same reasons. The expert systems accessed through electronic communications have the advantage of being far more accessible to group criticism and incremental improvement, and that is really their important function. They provide a medium for social wisdom and intercommunication of critical judgment to be compiled and put together in a consistent way. Of course, you can get majority rules that are total flops, but at least you can have a more accurate description of the consensus.

Current research and development in this field is concerned with more efficient ways of accomplishing knowledge extraction. This will require more intermediary languages so that experts can feel comfortable in expressing them-

selves in the habitual jargon of their own fields rather than through the somewhat stilted forms that most computer programs require. There are not many experts in knowledge engineering at the present time. There is little real science and much art in it, and only a few people have a history of practical experience. We need more meta-systems to enhance efficiency. As long as print publication dominates the field, we need systems that can read the print literature more effectively and economically. That need may, by itself, enhance still further the requirement that we have a more logically oriented language of expression. In other words, some stiltedness is desirable in order to reduce the ambiguities inherent in common language. We also need far more interoperability or compatability among expert systems that have been locally invented. Some people are working on mathematical expert systems; others on machine repair. There are people working on diagnostic systems, and others on chemistry. Such systems should be interoperable so that the expertise which is stored in one domain can be readily used in others without being painfully retranslated. There is no overarching framework to encourage that, although I hope that the appointment of Donald Lindberg as Director of the National Library of Medicine foretells how NLM will become the focal point for that interest and expertise.

The Defense Advanced Projects Research Agency has been the main sponsor of research on expert systems. Military decision-making poses requirements on designs for computer-aided support that even exceed those in medicine, and should have spillover advantages for our needs. A recent study of these efforts gives some perspectives on current and future technology. With reference to the new generations of computers that have been emerging in the 1980s, there are three varieties. On the one hand there are the huge mainframes or supercomputers that can perform hundreds of millions of operations per second, have vast memory, and are very well suited for prediction of the weather, the calculation of the aerodynamics of an air foil, or other areas of science which require very large numerical computations. Would that most biomedical science had a base sufficiently rigorous that we had wide ranging uses of that kind of super computation!

Quite different from those supercomputers are the novel architectures which are based on the fact that it is now possible to mass produce data processing chips. These processors are modestly fast at a few hundred thousand operations a second; the real costs are in wiring them up to one another. One could get the raw processing power of a supercomputer for only a few thousand dollars worth of chips, if we just knew how to take advantage of a very high order of parallelism. The main substantiation of that possibility is human intellectual activity, since there is no other way the human brain could be functioning. Our brains have very low order devices; synaptic switching times are in the order of a millisecond rather than a nanosecond, but there are trillions of them, and somehow those are wired up to enable us to perform certain kinds of computations at speeds that match what any of the supercomputers can do today.

Thirdly, there is the format of one chip per office, or the personal computer, where a chip has a terminal connected to it, a local memory, and a modem that enables it to be wired into the rest of the community. There are a few million of these around the country right now and collectively they constitute another kind of world brain of vastly different dimensions from anything that existed ten years ago. Another use of cheap chips is in all of the embedded computers that are generally not wired from one place to another, but which run individual instruments.

The medical world stands somewhat intermediate between the lab world demonstration and the real world life or death military arena. Ten million items in the database, and 10,000 relevant rules, characterize battle management. The flag officer on an aircraft carrier must defend his battle group against a wide variety of threats coming in from every quarter at very high rates. There are hostile forces and decoys, and questions about how many submarines are in the environment, how many planes are left to deploy, how rapidly the sorties for defense can be mounted, and so forth. The number of information items that have to be handled and processed is already more than human systems can handle at the present time; so mechanization is imperative.

Characteristic of both medical and military systems, the real cost of development is in the building of the expertise, not in its use. A thousand dollars per rule is a very conservative estimate of the cost of building an expert system. If unpaid medical students participate in the development, some of the true economic costs may be concealed, but this gives some notion of the costs of knowledge extraction and authentication.

Finally, what are the main uses of these new systems likely to be? Consider the integrity of the medical knowledge base on which the life or death of your patients depends, the integrity of your rules to reflect a wise consensus about the optimum form of therapy which incorporates in a consistent way all of the available information. I guess I am saying that I don't know how to do a human audit of the expert systems that have been brought together without an expert system to assist me in that process.

There is a deeper and more pervasive issue of how we are basing our entire economy, you might say our entire society, on rule systems that are formulated and managed by highly fragmented sets of experts such as the programmers in the banks, the people who run the credit systems, the people in the Internal Revenue Service. In a way, this system is essentially out of control. These people write the software, they deposit it deep in the bowels of the systems, they've never totally understood it, (nobody else ever does) and this is now done in languages that are not subject to any consistent form of authentication and audit. Everyone of us has had a disaster such as a phone bill with no bearing on anything that we had actually done. The question arises as to whether you can rely on systems that have that kind of fallibility, not in the hardware but in the interaction of the knowledge that is put into the system, when as physicians you have the responsibility for managing the welfare of your patients. So I think we need the meta-systems for the editing of the expertise so that human

intellect has some supportive tools in order to certify the integrity and authenticity of the expert systems we might rely on. Perhaps this conference will begin to get some consensual judgment on the important issues in this field for medicine. They apply to socially validated expertise in any complex endeavor, regardless of the use of the computers and communications that most relentlessly expose that expertise.

REFERENCES

Eisenstein, E. *The Printing Press as an Agent of Change*. Cambridge: Cambridge University Press, 1979

Gevarter, W. B. *Intelligent Machines*. Englewood Cliffs, NJ: Prentice-Hall, 1985.

——— *Report of the Defense Science Board Task Force on Military Applications of New-Generation Computing Technologies*. Washington, D.C.: Department of Defense, December 1984. (Defense Technical Information Center Report ADA152154).